REMARKS

Claim rejections under 35 USC 102

Claims 1-18 have been rejected under 35 USC 102(e) as being anticipated by Weber (6,631,448). Applicant respectfully traverses this rejection. However, Applicant first would like to mention to the Examiner that should he not agree with the arguments that are presented below, that he review the specification to locate subject matter that would render the patent application patentable to the satisfaction of the Examiner. That is, Applicant believes that examination of this patent application up to this point may have been unnecessarily antagonistic, and would like to make future examination more cooperative between Applicant and the Examiner. As such, if the Examiner can identify the subject matter within the patent application that he believes is patentable and call Applicant's representative, Mike Dryja, at the phone number listed below, then Applicant and the Examiner can work together quickly to reach a position at which the patent application is in a state for allowance. This would save significant time and expense on the part of both the Examiner and Applicant, and would further remove the present patent application from the dockets of the Examiner and Applicant.

With all the above stated, Applicant respectfully submits that Weber does not anticipate the claims. Claim 1 is discussed as representative of all the pending claims in this respect, insofar as the following limitation is concerned. This limitation is that the logic determines whether *a cache miss* relating to a memory unit should be transmitted to a sub-plurality of nodes, where this sub-plurality of nodes are in number less than the plurality of nodes but greater than one. That is, when you have a cache miss in the claimed invention, this cache miss is transmitted to a sub-plurality of nodes – not all the nodes, and not just one of the nodes.

Now, the Examiner has indicated on page 3 of the office action in particular that column 5, lines 55, through column 6, line 3 of Weber recite transmitting a cache miss to a number of nodes lesser than the all the nodes but greater than one node. Applicant respectfully disagrees here. This is why. If you review this portion of Weber carefully, you can see that it actually does

not relate to cache misses at all, but rather to cache line invalidation. That is, the issue is this: if I am a home node, and thus own a given cache line, and need to *invalidate* that cache line as stored in each of a number of other nodes (which in some situations will be more than one, but less than all the number of nodes), I send an *invalidate* instruction to these other nodes so that they know that the copy of this line that they have in their caches is no longer invalid. However, cache line *invalidation* is not the same as cache *misses*, which really does not have much to do with the home node of a memory line. That is, if I own a given line, I as the home node will really never in fact have a cache miss in that I will really never have to send the request to other nodes – because I own the line and do not have to send the request to other nodes.[1]

Rather, Applicant submits that the relevant portion of Weber as to what you do in cache misses is column 6, lines 11-34. This is the only part of Weber that mentions cache misses – as compared to the excerpt discussed in the previous paragraph – and therefore is the part of Weber

[1] This excerpt of Weber reads as follows:

> The home node directory 1960 (140) tracks those nodes 200 that have requested a read-only copy 1923 of the cache line 1900 so that when a node wants to update a line 1900, the directory 1960 (140) knows which nodes to selectively invalidate. In addition, the directory 1960 (140) tracks the node, if any, that owns the memory line 1900. The home node 1920 knows all the remote nodes 1950 within the shared-memory site 120 that have requested a read-only copy of the home node's 1920 memory line 1900, and also the remote node 1930 that has requested write access to the memory line 1900. When home node 1920 must communicate with remote nodes 1930, 1950, it does so selectively based upon which nodes have accessed the line rather than to all nodes within site 120. Thus, the directory-based cache coherence protocol achieves coherence by point-to-point communication between nodes 100 rather than broadcast invalidation.

Thus, this portion of Weber talks about what you do in invalidation, and/or what you do when a home node has to communicate with nodes that are caching its owned memory lines. This portion of Weber does not talk about cache misses, which is what the claim limitation in question talks about. Therefore, in essence, this part of Weber is basically irrelevant.

that has to recite the claim limitation in question for Weber to anticipate the claimed invention. Now, if you read this part of Weber carefully, what happens is that if you are a node that has a cache miss, you issue a request onto your local P6 memory bus, and if the memory line in question resides on a different node, then the node issues "a network request for the cache line 1900 across Interconnect 110 ***to home node 1920***." The home node then figures out what to do – i.e., if the memory line is locally cached, if it has been updated in a remote cache, etc., and reports back to the original node issuing the request.[2]

Now, compare what Weber does with a cache miss as compared to the claimed invention. In the claimed invention, you transmit the cache miss to a number of nodes *greater than one*, but not to all the nodes. So if you have four nodes, you do not transmit the cache miss to all four nodes, nor do you transmit the cache miss to just one node. You would transmit the cache miss

---

[2] This excerpt of Weber reads as follows:

> If processor 1941 does not find the data in cache 1942 (a "cache miss"), the processor issues a request onto P6 memory bus 230 for the data. The bus request includes the real address of memory line 1900 and the identification of home node 1920 for memory line 1900. **If memory line 1900 resides on a different node, MCU 1943 (130) of local node 1940 generates a network request for the cache line 1900 across Interconnect 110 to home node 1920.** When the MCU at home node 1911 (130) receives the network request, it stages the request onto home node's memory bus 1912 (230). The network request then snoops bus 1912 (230) to determine if memory line 1900 is cached locally 1921. If not cached, the network request obtains control of memory line 1900 in the node's memory 1910 (240). Thus, the action of MCU 1911 (230) at home node 1920 is like the action of a local processor 210 accessing memory line 1900 within node 200. Depending on the status of memory line 1900, MCU 1911 (230) will take the appropriate action to give the requested memory line 1900 to local node 1940.

Thus, for a cache miss in Weber, you only transmit the cache miss (i.e., the network request) to the home node, which, of course, is just *one node*. By comparison, in the claimed invention, you transmit the cache miss to a number of nodes that is *greater than one* but less than all the nodes, as is discussed in more detail later in this office action response.

to two or three nodes. However, Weber just transmits a cache miss (i.e., a "network request") to the *home node* for the memory line in question. That is, Weber transmits a cache miss to just one node. Therefore, Weber cannot anticipate the claimed invention, because it does not transmit a cache miss to a sub-plurality of nodes that is lesser in number than all the nodes but greater than one, to which the claims are in essence all limited.

Applicant wants to discuss this a little further now, to make this point even more clear. Applicant concentrates on two sentences in Weber. First, "the directory-based cache coherence protocol achieves coherence by point-to-point communication between nodes." (Col. 6, ll. 3-5) Second, "if processor 1941 does not find the data in cache 1942 (a "cache miss")" and "[i]f memory line 1900 resides on a different node, MCU 1943 (130) of local node 1940 generates a network request for the cache line 1900 across Interconnect 110 to home node 1920." (Col. 6, ll. 9-19) The latter sentence is what has been discussed in the previous paragraph; the former sentence says that Weber uses point-to-point communication.

Now, Applicant wants to underscore the point that the patent application as filed discussed this type of cache coherency system, in which cache misses are sent to just one node, in a point-to-point manner, as being prior art, within the background section of the specification, as follows:

> In a unicast, or point-to-point or directory-based, cache coherence protocol, the originating node always sends a single request – i.e., the cache miss – for the contents of the memory unit to one other node. Where the memory unit is local to the originating node, such that the originating node is the home node for the memory unit, the originating node sends a single request for the contents of the memory unit to the remote node that has modified the contents of the memory unit. In response, the remote node sends the contents of the memory unit, as have been modified, back to the originating node. Where the memory unit is remote to the originating node, the originating node sends a single request for the contents of the memory unit to the home node for the memory unit. Because the home node for the memory unit may not actually hold the current contents of the memory unit, it may have to forward the request to a third node, which may have modified the contents of the memory unit.

> Unicast, or point-to-point or directory-based, cache coherence protocols minimize total communication traffic among the nodes of a multiple-node shared-memory system, because cache-coherence requests resulting from cache misses are only sent from an originating node to one or two other nodes, in the most frequent scenarios. Therefore, such systems generally have good scalability, because adding nodes does not add an inordinate amount of communication traffic among the nodes for cache coherence purposes. However, latency may suffer within systems using such cache coherence protocols, since the recipient node of the originating node's request may not actually have the current contents of the desired memory unit, requiring the recipient node to forward the request to another node.

(P. 4, ll. 1-22) Thus, the claimed invention is inherently directed to *something else*, specifically a methodology in which you send a cache miss to more than one node but not to all the nodes.

In other words, Applicant refers the Examiner to the summary section of the specification of the patent application as filed, which explains why the claimed invention is advantageous over such point-to-point, or unicast, methodologies in which a cache miss is sent to just one other node (such as the home node, as in Weber), as follows:

> Embodiments of the invention provide for advantages over the prior art. Whenever possible, logic is used to determine when broadcasting a cache miss to all the nodes of a system is not necessary to ideally reach the owning node of a memory unit without reissuing the cache miss, such that selective broadcasting suffices to ideally reach the owning node of the memory unit without reissuing the cache miss. Thus, embodiments of the invention are advantageous over unicast-only protocols that always unicast cache misses, because unicast-only protocols will necessarily incur forwarding latency in at least some instances, which is at least substantially avoided by embodiments of the invention.

(P. 8, l. 16, through p. 9, l. 2) Thus, the claimed invention does not use point-to-point, or unicast, transmission of cache misses, unlike Weber. Hopefully the foregoing discussion makes it clear how Weber differs from the claimed invention in this respect – if not, though, Applicant's representative, Mike Dryja, is very much amenable to being called by the Examiner at the phone number listed below, so that Mike can explain this in more detail. The take home point here is that in Weber, you transmit a cache miss from the missing node to just one node, whereas in the claimed invention, you transmit it to more than one node (but not to all nodes).

Applicant would like to conclude this discussion with another point that should also underscore the difference between the invention being considered here and Weber. Now, the background section of the patent application as filed talks about why point-to-point cache miss methodologies can incur undesired latency, because when you transmit a cache miss to the home node, the home node may have to reissue this cache miss to one or more other nodes, and so on. So it may take a while before the responsible node actually receives the cache miss. The claimed invention's solution to this problem is to say, "OK, we are going to transmit the cache miss to a selective group of nodes – more than one in number, but less in number than all the nodes." The idea is that you avoid latency, because the group you select is likely to include the responsible node for the memory line that is the subject of the cache miss. (The other idea is that you avoid unnecessarily putting too much traffic on the interconnect by forwarding the cache miss to *all* the nodes, but this point is not as important here in light of Weber, of course.)

Now, interestingly, Weber achieves a *different solution* to the problem outlined in the background section of the patent application as filed. Basically, Weber says that, "OK, we are going to transmit the cache miss to just one node – and if that node is not responsible for the memory line in question, we are going to *make* this node responsible." As such, unicast-oriented latency is also avoided by Weber, but in a different way as compared to the claimed invention. Whereas the claimed invention reduces latency by transmitting a cache miss to more than one node, Weber reduces latency by transmitting the cache miss to just one node – but then thereafter ensuring that the cache miss does not have to be transmitted (i.e., reissued) to any other node from that one node.[3] The point that Applicant wants to make here is that Weber seems to solve

[3] For instance, once the home node receives the transmitted cache miss/network request, this is what it does in Weber.

> When the MCU at home node 1911 (130) receives the network request, it stages the request onto home node's memory bus 1912 (230). The network request then snoops bus 1912 (230) to determine if memory line 1900 is cached locally 1921. If

the same problem that the invention does – but Weber solves it in a different way! Because Weber teaches a cache miss methodology different than that of the claimed invention, in other words, Weber cannot anticipate the claimed invention.

Conclusion

Applicants have made a diligent effort to place the pending claims in condition for allowance, and request that they so be allowed. However, should there remain unresolved issues that require adverse action, it is respectfully requested that the Examiner telephone Mike Dryja,

---

> not cached, the network request obtains control of memory line 1900 in the node's memory 1910 (240).

(Col. 6, ll. 19-25) Thus, there are at least three different approaches you can do. First, you can follow the prior art explained in the background section of the patent application as filed, in which a cache miss is transmitted to the home node, which then may have to retransmit the cache miss to another node that is actually responsible for the memory line in question. This approach is disadvantageous, due to latency incurred in the "double hop" – from the cache miss node to the home node and then to the home node to the responsible node. Second, you could do what Weber does, where a cache miss is transmitted to the home node, and if the home node is not the responsible node, it takes responsibility (i.e., control) for the memory line in question anyway. Weber's approach avoids latency, because you only have a "single hop" from the cache miss node to the home node. Third, you could do what the claimed invention does, where a cache miss is transmitted to *more than one node* but not to all the nodes. Presumably, one of the nodes to which the cache miss is transmitted will be the responsible node. The invention's approach also avoids latency, because you only have a "single hop" from the cache miss node to the home node – but the invention's approach is nevertheless different than the approach that has been claimed, because Weber transmits the cache miss to just one node – the home node – whereas the invention transmits the cache miss to *more than one node*.

Applicant's representative, at 425-427-5094, so that such issues may be resolved as expeditiously as possible. Applicant again wants to underscore that Applicant would like for the examination process to be cooperative between the Examiner and Applicant, and not antagonistic or adversarial. To that end, if the Examiner does not understand Applicant's discussion above – which Applicant would completely understand, since Applicant knows that the application is very technical, and Applicant sympathizes that the Examiner did not have the benefit of having drafted the patent application like Applicant's representative, Mike Dryja, did – or if the Examiner understands Applicant's discussion but believes that one or more other limitations have to be added to the claims to render them allowable, the Examiner is very strongly encouraged to contact Mike Dryja. That is, Applicant would like to move the present patent application to allowance as quickly as possible, and requests that the Examiner propose claim amendments that he believes would render the claims allowable. Applicant is very amenable in this respect to consideration of such proposals by the Examiner. In any case, for these reasons, this application is now considered to be in condition for allowance and such action is earnestly solicited.

Respectfully Submitted,

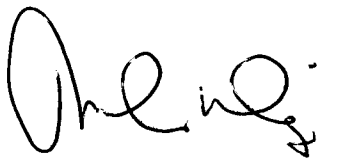

August 5, 2007
Date

Michael A. Dryja, Reg. No. 39,662
Attorney/Agent for Applicant(s)

Law Offices of Michael Dryja
1474 N Cooper Rd #105-248
Gilbert, AZ 85233
tel: 425-427-5094
fax: 206-374-2819